



ELSEVIER

Contents lists available at ScienceDirect

Journal of Housing Economics

journal homepage: www.elsevier.com/locate/jhec

Re-examining the impact of housing wealth and stock wealth on retail sales: Does persistence in wealth changes matter? ☆



Richard Ashley, Guo Li *

Department of Economics, Virginia Tech, United States

ARTICLE INFO

Article history:

Received 8 January 2014

Revised 11 September 2014

Accepted 19 September 2014

Available online 2 October 2014

JEL Codes:

E21

C22

Keywords:

Wealth effect

Frequency dependence

Persistence dependence

ABSTRACT

Case, Quigley and Shiller (2005, 2013) quantified stock versus housing wealth effects on quarterly state-level retail sales, which they interpret as an approximate measure of household consumption spending. We investigate the variation of these wealth effects with the persistence of each kind of wealth fluctuation in an estimated linear dynamic fixed-effects model, allowing for both cointegration and endogeneity. Retail sales respond most strongly to housing wealth fluctuations which persist for one to four years, whereas the response to stock wealth fluctuations is smaller and is concentrated on fluctuations with a persistence of either less than a year or more than four years. These differential persistence effects point to a need for a richer theoretical formulation in this area.

Published by Elsevier Inc.

1. Introduction

A substantial literature has arisen that compares the wealth effect due to housing wealth fluctuations with the wealth effect due to financial wealth fluctuations (E.g., see Edison and Sløk (2001), Case et al. (2005) and Shirvani and Wilbratte (2011)). This issue is important because both of these kinds of wealth fluctuations have played major (albeit likely intertwined) roles in triggering and/or extending major macroeconomic episodes in the last few decades and because a fluctuation in each of these household wealth variables calls out for a different set of preventive and/or reactive government policies. Some studies argue that the transitory nature of the changes in stock prices causes them to have a

smaller impact on consumption than changes of similar size in the value of other assets – e.g., Benjamin et al. (2004); Case et al. (2013). However, Dvornak and Kohler (2007) find opposite results in modeling the seven provinces of the Australian economy, where fluctuations in financial wealth appear to have larger impacts than fluctuations in housing wealth. Belsky (2010) found similar consumption effects from real estate and corporate equity fluctuations, both at a magnitude of 5.5 cents on the dollar. Similarly, Carroll et al. (2011) find that the financial wealth effects grows to be more than four-to-ten cents on the dollar over the years following a shock. Their results also suggested that about 80% of the housing wealth effect is realized in one year, whereas a long run effect from the stock market takes five years to approach 80%. Engelhardt (1996) finds asymmetric wealth effects: changes in consumption are significantly associated only with drops in housing values. Case et al. (2005, 2013) used state-level panel data on retail sales, household financial wealth, and

☆ We are grateful to John Leahy, Randal Verbrugge, and Kwok Ping Tsang for helpful comments on an earlier draft.

* Corresponding author.

household housing wealth to analyze the differential impacts of these two kinds of household wealth on household consumption, the latter of which is in these studies (and here) approximated at the state level by total retail sales.¹

Considering the contradictions in the literature about the different types of wealth effect, the main contributions of the present study are to extend the wealth effect analysis in two ways. First, we estimate a single model for the Case et al. (2005, 2013) data using a standard dynamic panel framework, including both of the two kinds of household wealth effects, and also simultaneously allowing for both endogeneity in the two household wealth growth rates (i.e., financial and housing) and also allowing for cointegration amongst the concomitant key levels variables.² This allows us to compare the relative consumption impact of financial wealth fluctuations versus housing wealth fluctuations in a single model accounting all at once for all of the major econometric features of their data set. Our results are therefore both more economically coherent and statistically more valid than those obtained previously.

Second, and most crucially, we use a recently developed econometric method, due to Ashley and Verbrugge (2006, 2009), and Ashley et al. (2014) – which enables us to examine how the relationship between the growth rate of retail sales varies with the *persistence* of each of these kinds of household wealth fluctuations. This method employs one-sided Fourier filtering to decompose each of the two wealth growth rates into several persistence level components; these components are in each case constructed so as to add up to the original wealth series. This decomposition is described in Section 3 and in Appendix C; it is there contrasted to alternatives, such as HP filtering.

Why is it important to analyze the manner in which the consumption impact of fluctuations in these two household wealth varies with the persistence level of the fluctuations? First, note that the ‘persistence’ referred to here is the persistence of the recent fluctuations in the stationary – i.e., $I(0)$ – growth rates of the these two household wealth time series. For example, a positive fluctuation in the growth rate of housing wealth in a particular state which is a part of recent pattern of positive fluctuations is what we are terming a persistent fluctuation in this kind of wealth, whereas an isolated positive fluctuation in housing wealth is what we are referring to as a less persistent fluctuation. Since we find that each of the two wealth effects does in fact depend on the persistence level of the wealth fluctuation, any model which does not account for this persistence-variation is providing a single, inconsistent

estimate of the wealth effect, averaged over all persistent levels.³

In addition – beyond simply finding persistence dependence in both the stock wealth and housing wealth coefficients – the dependence on persistence level which we find in these coefficients is economically interesting in form. For one thing, this dependence is quite different for each of these two kinds of household wealth. For another, our results with regard to the form of the persistence dependence are in each case a bit surprising. In particular, while there is no existing theory available to predict this form, an informal appeal to the Permanent Income Hypothesis would at least suggest that these wealth effects on consumption would both be monotonically increasing in the persistence level of the wealth fluctuations. Such is not the case, as discussed below.

We note that this is fundamentally an empirical paper: we do not provide a new theory of household consumption predicting the pattern of wealth-persistence effects which we find. Rather, we hope that our intriguing empirical results will motivate the development of such theories. It is reasonable to speculate in that direction, however, and we do so in our concluding section.

The rest of this paper is organized as following: Section 2 summarizes the Case et al. (2013) data; Section 3 briefly discusses our econometric terminology and empirical model; Section 4 presents the results, and Section 5 concludes.

2. Data

We use state-level per capita owner-occupied housing wealth, per capita financial wealth and per capita household consumption, as imputed in Case et al. (2005, 2013). This is virtually the only data set that has both the financial wealth and housing wealth disaggregated to the state level; the imputation covers a significant period of time, from the first quarter of 1978 to the fourth quarter of 2012. This data set offers several advantages for our persistence decomposition analysis: (1) The increase in both forms of wealth has been quite unequally distributed across geographic units; this panel offers the advantage that the variable definitions are uniform across different states. (2) This data set also makes it possible to define an error correction term based on the relationships between the level variables for each state. (3) The sample spans over thirty years of US economic history, with a total of 135 quarterly observations per state. This long panel allows us to easily specify windows 16 quarters in length for the Fourier analysis, so that only fluctuations with a

¹ Despite omitting important components, such as services, etc., retail sales have been argued by Case et al. (2005, 2013) and Elbourne (2008) to be a reasonable proxy for household consumption spending at the state level. In particular, Elbourne (2008) finds a sample correlation of 0.95 between retail sales and consumption at the national level.

² The models in Case et al. (2005, 2013) are not able to deal with all of these econometric features at once in a single model.

³ Our analysis is thus analogous to estimating the marginal propensity (MPC) out of ‘temporary’ income and comparing this to an estimate of the MPC out of ‘permanent’ income – as in Permanent Income Hypothesis analyses, such as Campbell and Mankiw (1990) – rather than only estimating an average of these two MPC values. Previous work in the wealth impact area – per this analogy – has only analyzed the average MPC, comparing the average MPC out of financial wealth to the average MPC out of housing wealth, ignoring the fact that – as we uncover – the consumption impact of a fluctuation in either kind of wealth depends on the persistence of this fluctuation.

Table 1

Summary statistics for housing wealth, stock wealth, income and retail sales (growth rates).

Growth rates	Observations	Mean	Std. Dev.	Min	Max
<i>Pre-crisis periods (1978Q1–2007Q4)</i>					
Housing Wealth	6681	0.004	0.034	−0.482	0.644
Stock Wealth	6681	0.013	0.083	−2.646	2.672
Retail Sales	6681	0.003	0.017	−0.297	0.107
Income	6681	0.004	0.015	−0.182	0.185
<i>Financial crisis and afterwards (2008Q1–2012Q4)</i>					
Housing Wealth	918	−0.012	0.034	−0.145	0.103
Stock Wealth	918	−0.006	0.083	−0.735	0.291
Retail Sales	918	−0.003	0.021	−0.084	0.043
Income	918	−0.001	0.016	−0.067	0.073

persistence level more than four years in length are aggregated into the lowest frequency component.

One disadvantage of this dataset, however, is that per capita consumption is approximated at the state level by total retail sales. This is the same approximation used in Case et al. (2005, 2013), but it is not unique to them. Moreover, per footNote 1, consumption and restail sales are highly correlated at the national level.

We also note that Case et al. (2005, 2013) restricts the growth rate in household financial wealth solely to the growth rate in households' holdings of mutual funds due to data availability. To the extent that the sample fluctuations in the growth rate of household stock holdings and the growth rate of other financial assets (e.g., bonds) are uncorrelated, this implicit variable omission from the regression model is econometrically inconsequential; to the (more likely) extent that these fluctuations are quite highly correlated, then one can interpret the stock wealth growth rate fluctuations in this data set as a reasonable proxy for the concomitant fluctuations in the growth rate of total household financial wealth.

Table 1 provides summary statistics for the Case et al. (2013) data prior to decomposition into components of differing persistence levels. During the 1978Q1–2007Q4 pre-crisis sample period, the growth rate of per capita housing wealth averaged 0.4% and the stock wealth growth rates averaged 1.3%. During and subsequent to the financial crisis period, however, both of these wealth growth rates exhibit a downward trend, with growth rates averaging −1.2% for housing wealth and −0.6% for stock wealth. The growth rate of retail sales also exhibits a sharp contraction during the financial crisis: it grew steadily with a rate of 0.3% per year before the onset of the financial crisis, while decreasing at a rate of 0.3% during the financial crisis. Similarly, income grew at 0.4% prior to the financial crisis, and shrank at a 0.1% rate thereafter. Table 1 also displays the fact that there is more sample variation in the growth rate of stock wealth than of housing wealth.

The decomposition method we use here originated in Ashley and Verbrugge (2006) and Ashley and Verbrugge (2009); it was further developed in Ashley and Tsang (2012) and Ashley et al. (2014). The latter two papers describe the method in detail, so here we provide only a summary, in Appendix A.

Previewing the results of this decomposition, Fig. 1 displays a time plot of the resulting sample data for the low, middle, and high frequency components of the growth rate

in housing wealth and stock wealth for a typical state (New York); these add up, by construction, to the Case et al. (2013) data set.⁴ The decomposition of the growth rate of housing wealth is in the top portion of Fig. 1, whereas the decomposition for the growth rate of stock wealth is displayed in the bottom portion. The difference in the persistence levels of the three frequency components is clear from comparing the three lines in each plot: The dotted line represents the highest-frequency or least-persistent component of the wealth variables; it includes the fastest/briefest fluctuations in the time series. The dashed line represents the mid-frequency component; here fluctuation reversal is less prompt. The solid line shows the lowest frequency components, which captures the nonlinear trend, and also any mean-reversal behavior with a period longer than four years in length. Comparing the low frequency bands of the growth of housing wealth to the low frequency bands of the growth of stock wealth, we see that the New York stock wealth time series is more volatile – i.e., less persistent – than the housing wealth series.

3. Empirical model

Our empirical model is specified as the following:

$$\Delta s_{it} = \lambda'_i \Delta x_{it} + (\beta^{-stock})' \begin{bmatrix} \Delta a_{it}^{stock_low} \\ \Delta a_{it}^{stock_mid} \\ \Delta a_{it}^{stock_high} \end{bmatrix} + (\beta^{-house})' \begin{bmatrix} \Delta a_{it}^{house_low} \\ \Delta a_{it}^{house_mid} \\ \Delta a_{it}^{house_high} \end{bmatrix} + \varepsilon_{it} \quad (1)$$

For the $i = 1 \dots 51$ states (including the District of Columbia) and $t = 1 \dots 135$ quarters of data assembled by Case et al. (2013), Δx_{it} includes the other explanatory variables they specify: e.g., the intercept, $\Delta s_{i,t-1}$, and the error-correction term. The vector-valued parameter λ'_i is notated with subscript i to allow for its intercept component to be state-specific in the usual fixed-effects fashion.

⁴ Windows-based software is available from the authors, which converts a file containing a panel of sample data into a spreadsheet file whose columns contain the corresponding frequency components for model parameter estimation. This program allows one to specify the window length (M) and partitions the input data series into all M possible components, so one can make an alternative aggregation choice to that used here.

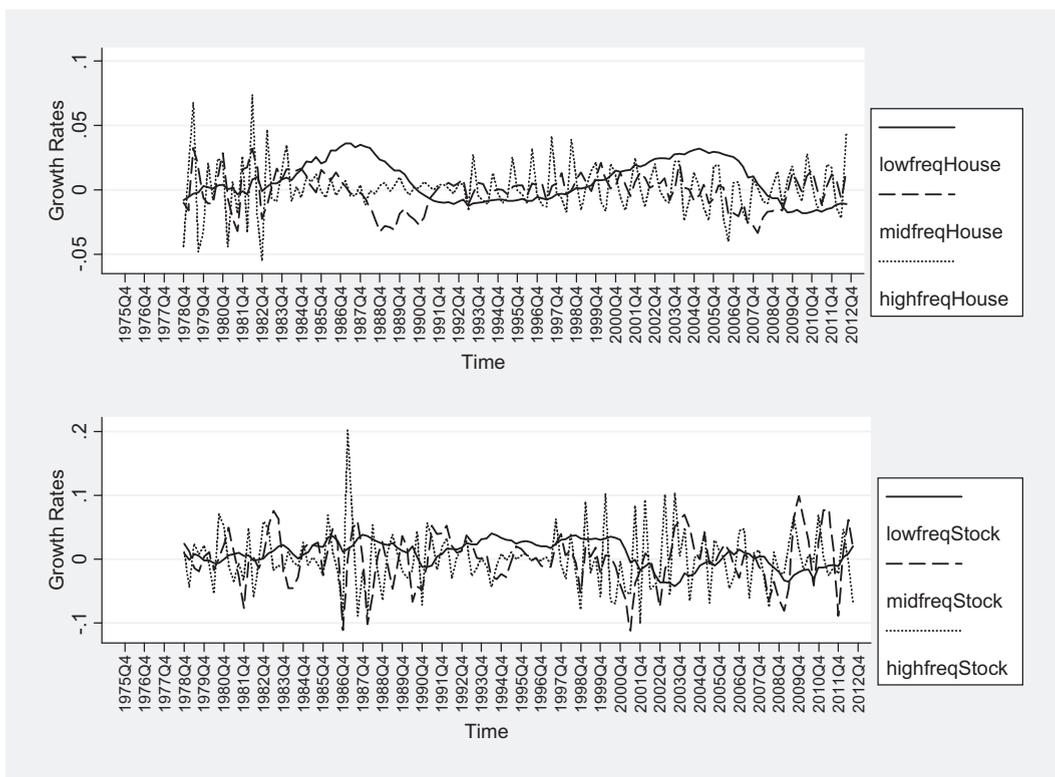


Fig. 1. Time plots of decomposed housing wealth and stock wealth growth rates for the New York State data.

Eq. 1 decomposes the two Case et al. (2013) wealth growth rate variables (Δa_{it}^{stock} and Δa_{it}^{house}) each into the three persistence levels indicated below:⁵

1. $\Delta a_{it}^{wealth_type_low}$ is defined as the component of $\Delta a_{it}^{wealth_type}$ which corresponds to fluctuations persisting for more than sixteen quarters (the “low-frequency” or “high-persistence” component).
2. $\Delta a_{it}^{wealth_type_mid}$ is defined as the component of $\Delta a_{it}^{wealth_type}$ which corresponds to fluctuations persisting for more than four, but less than sixteen quarters (the “mid-frequency component”).
3. $\Delta a_{it}^{wealth_type_high}$ is defined as the component of $\Delta a_{it}^{wealth_type}$ which corresponds to fluctuations persisting for four quarters or less (the “high-frequency” or “low-persistence” component).

Note that this decomposition actually *partitions* each wealth growth rate series into these three “frequency bands” or “persistence levels”. Thus, the three persistence components for each wealth series sum up precisely to the

⁵ Here “wealth_type” can be either housing wealth (housing) or financial wealth (stock). As is explained in Appendix A, the 16-quarter moving window used here – so as to make the filtering one-sided, and hence immune to statistical distortions due to any feedback in the consumption-wealth relationships – also limits the number of mathematically possible frequency components in the present study to just nine. We aggregate those nine components into these three wealth types primarily for ease of interpretation.

original Case et al. (2013) data series. Therefore, allowing for persistence dependence in the relationships between the retail sales growth rate and the two wealth growth rates reduces to simply replacing each wealth growth rate series by a weighted sum of its three persistence components, and then testing the null hypothesis that the three estimated weights – either $(\beta_{1i}^{stock}, \beta_{2i}^{stock}, \beta_{3i}^{stock})$ or $(\beta_{1i}^{house}, \beta_{2i}^{house}, \beta_{3i}^{house})$ – are equal.

A lagged error-correction term is included as one component of the vector Δx_{it} . The coefficient on the error-correction term in our model quantifies the effect on the current growth rate of retail sales due to the level of retail sales being, for example, above the value one would expect, given the long-term equilibrium relationship this level variable has with other cointegrating variables (such as the level of household income and of each kind of wealth). The specification of this term is described in Appendix B.

We also follow the Case et al. (2005, 2013) contemporaneous specification of the retail sales-wealth relationships and therefore allow for likely endogeneity in the six wealth growth rate components. Lastly, we separate the wealth effects estimates for the financial crisis period (2008Q1–2012Q4) from those of the prior periods, as there appears to be a structural shift at that point; this issue is discussed below in Section 4 and Appendix C. We estimate Eq. 1 using standard GMM fixed-effects dynamic panel data estimation methods, allowing for both a lagged dependent variable, the endogeneity in the wealth variables, a state-specific linear time trend, and – per fixed

effects – for non-homogeneity in the state-level model intercept.⁶

4. Results

Tables 2 and 3 display the coefficient estimation and inference results on the set of six persistence-specific wealth growth coefficients – three persistence levels on each of two wealth growth rates – for the sample period of 1978 to 2007Q4, prior to the onset of the financial crisis. Since the dependent variable in Eq. 1 is the growth rate in retail sales and each of the six wealth explanatory variables is a persistence-specific component of a growth rate, each coefficient in Table 2 is interpretable in the usual way, as an elasticity estimate. Thus, for example, Table 2 indicates that a 1% increase in housing wealth that is expected to persist for 1 to 4 years should on average, all else equal, lead to a 0.41% increase in retail sales.

Analogous results on the sample period 2008Q1 to 2012Q4 – i.e., including the data from the financial crisis through the end of the available sample period – are of lesser interest because it is evident that the Case, et al. (2013) model summarized in Eq. 1 breaks down at that point. In particular, we tested for this structural shift by including a coefficient-shift dummy variable (set to one for all time periods subsequent to 2007Q4) on every explanatory variable. The joint null hypothesis of coefficient stability is rejected with $P < 0.0005$, and this result is not sensitive to minor changes in the onset date. More importantly, several of the resulting wealth coefficient estimates are significant and negative in this later period; these results are tabulated in Appendix C.

Turning to Table 2, the first thing to notice is that the estimated coefficients on the six frequency/persistence-specific wealth growth rate variables are all positive – as one would expect – except for a very small (and statistically insignificant) negative coefficient on $HighfreqHouse_{i,t}$. Fluctuations in housing wealth which reverse within four quarters evidently have no impact on household consumption, quantified here by retail sales. Continuing:

- The consumption responses to fluctuations in both kinds of household wealth are substantially frequency/persistence-specific, with a quite different frequency pattern for each of the two wealth types.
- For both low and medium frequency wealth fluctuations, the consumption responses to fluctuations in housing wealth are significantly larger than those due to fluctuations in stock wealth, economically and statistically.

More specifically, the inference results in Table 3 indicate that the null hypothesis that the three frequency/persistence-specific coefficients on each household wealth fluctuation variable are equal can be rejected with $P < 0.0005$; the joint test that the four parameter restrictions corresponding to testing the null hypothesis that this is the case for both household wealth variables at once can

Table 2

Estimation results for wealth coefficients Eq. 1.

Lowfreq Stock	Midfreq Stock	Highfreq Stock	Lowfreq House	Midfreq House	Highfreq House
0.0937*** (0.0226)	0.0491** (0.0194)	0.119*** (0.0159)	0.190*** (0.0334)	0.410*** (0.0428)	–0.0399 (0.0557)

Note: *** $P < 0.01$, ** $P < 0.05$, * $P < 0.1$. Robust standard errors in parentheses.

Table 3

Inference results for wealth coefficients in Eq. 1.

<i>P-value for Ho: no frequency dependence</i>		
Stock wealth <0.0005	Housing wealth <0.0005	Both wealth variables <0.0005
<i>P-value for Ho: housing wealth effect equals stock wealth effect</i>		
Low frequency band 0.0038	Mid frequency band <0.0005	Both bands <0.0005

Note: “Low Frequency Band” corresponds to the wealth components whose fluctuations are most persistent; similarly, “High Frequency Band” corresponds to the least persistent fluctuations. Since the coefficient on $HighfreqHouse_{i,t}$ is statistically insignificant (with the opposite sign), the hypothesis that its coefficient equals that on $HighfreqStock_{i,t}$ is not relevant and thus not included in the joint test of the null hypothesis that the housing wealth effect equals the stock wealth effect.

be rejected with $P < 0.0005$ also. Substantial frequency dependence is clearly an important feature in both parts of the household consumption-wealth relationship.

Moreover, the pattern of the frequency dependence is quite distinct for fluctuations in the two kinds of household wealth. For housing wealth, it is fluctuations in the mid-frequency band (i.e., $MidfreqHouse_{i,t}$, which corresponds to fluctuation with periods between five quarters and four years) which are overwhelmingly most important; in fact – as noted above – fluctuations in $HighfreqHouse_{i,t}$ appear to have no statistically significant impact on consumption at all.⁷ Notably, the coefficient on $MidfreqHouse_{i,t}$ substantially (and significantly) exceeds that on $LowfreqHouse_{i,t}$; this is not what one would expect. In contrast, for household stock wealth, it is clearly fluctuations in $LowfreqStock_{i,t}$ and $HighfreqStock_{i,t}$ which influence consumption the most.

Next we compare the relative impact of a fluctuation in each kind of wealth, disaggregated by persistence level. For household wealth fluctuations in the mid-frequency band, the estimated coefficient on the growth rate in housing wealth substantially exceeds that on stock wealth; the null hypothesis that the difference in these two coefficients is zero can be very strongly rejected, with $P < 0.0005$. Similarly, the null hypothesis of equal housing and stock wealth impacts can also be strongly rejected for the low-frequency band ($P = 0.006$). In contrast, for high frequency wealth fluctuations, the impact on consumption of changes in housing wealth is statistically insignificant, whereas an increase in stock wealth has a modest positive consumption impact which is very statistically significant: the null

⁶ Models including quadratic trends or no trend at all yielded higher BIC values and similar results; these are available from the authors on request.

⁷ Because the five frequency components aggregated into $HighfreqHouse_{i,t}$ are individually so noisy, several of these components yield perverse (negative) coefficient estimates when included individually. However, their sum is not significantly different from zero at the 1% level; in view of the number of tests being individually examined here, we interpret those estimates as random sampling variation.

Table C.1

Estimation results for Eq. 1: non-wealth variables in the pre-crisis periods.

Lag (1) Dependent	Lag (2) Dependent	Error correction	Change in
Variable	Variable	Term	Income
−0.0449** (0.0175)	0.104*** (0.0163)	−0.121*** (0.0150)	0.609*** (0.0648)

Note: *** $P < 0.01$, ** $P < 0.05$, * $P < 0.1$. Robust standard errors in parentheses.

hypothesis that this coefficient is zero can be rejected with $P < 0.0005$. The inference results in Tables 2 and 3 thus allow us to analyze this issue more informatively than was possible in previous studies, because we address it separately for wealth fluctuations at each of the three frequency or persistence levels⁸ (see Table C.1).

Since the dependent variable in Eq. 1 is the growth rate in retail sales and each of the six wealth explanatory variables is a persistence-specific component of a growth rate component, each coefficient in Table 2 is interpretable in the usual way, as an elasticity estimate. Thus, for example, Table 2 indicates that a 1% increase in housing wealth that is expected to persist for 1 to 4 years should on average, all else equal, lead to a 0.41% increase in retail sales.

5. Conclusions

This paper is the first effort in the literature to empirically examine the differential impact of both housing and stock wealth fluctuations on state-level retail sales across different levels of persistence in these wealth fluctuations. As noted earlier, we follow Elbourne (2008) and Case et al. (2005, 2013) in interpreting these results as shedding light regarding the analogous impacts of fluctuations in these two kinds of household wealth on household consumption spending. While we leave the formal construction of new theories of household consumption spending based on these results to others – whose comparative advantages are better aligned to that task – we view our empirical findings as new ‘stylized facts’ which any such theories should be able to generate, but we do conjecture below on what these new “facts” might mean.

One clear-cut finding that emerges from our empirical work is that, at least for the period prior to the 2008 financial crisis, the persistence level of wealth fluctuations is highly consequential to their impact on household consumption spending.

Thus, for example, our results confirm the Case et al. (2005, 2013) result that state-level fluctuations in household consumption depend much more strongly on fluctuations in housing wealth than on fluctuations in stock

wealth; but we find that this is the case only for wealth fluctuations with a reversion period of more than a year. These results reinforce the notion that housing is more than just a financial asset: housing services are part of consumption, in addition to housing itself playing a role as collateral in financial markets.

We also find that households typically ignore short term, quickly-reversed shocks to housing wealth, but respond much more intensely if the shock lasts longer than a year. Our result that housing wealth fluctuations appear to have no statistically significant effect on household consumption at all for reversion periods of a year or less could be due to the relative illiquidity of housing wealth. But this still leaves our finding that a low frequency fluctuation in the housing wealth growth rate has a larger consumption impact than a medium frequency fluctuation of equal size something of a puzzle: one would expect the opposite.

In contrast to these results on housing wealth, we find that fluctuations with a reversion period of between one and four years in wealth held as stock appear to have no statistically significant effect on household consumption, but have a statistically significant, albeit modest, positive impact on household consumption spending when these fluctuations are of either low persistence (i.e., tend to revert within a year) or very high persistence (i.e., persist for more than four years). As with the analogous housing wealth result, it is intriguing to find that the consumption impact of a stock wealth fluctuation is not monotonically increasing in the persistence level of the fluctuation. This result suggests to us that financial wealth impacts household consumption spending through two distinct economic mechanisms. We conjecture that one of these mechanisms operates through fluctuations in a stochastic discount rate whereas the other operates through an asset pricing channel.

In closing, our results clearly indicate that the dynamics by which household consumption spending – at least as proxied for by retail sales – reacts to fluctuations in wealth held as stock (versus fluctuations in wealth held as housing) are distinctly different. In particular, both the overall size and the dynamics of these reactions depend on the persistence of each kind of wealth fluctuation. Moreover – interestingly – our results clearly indicate that more “permanent” fluctuations in either kind of wealth do not necessarily have a larger impact on consumption spending. These strictly empirical results suggest to us the existence of a rich vein for theorists to mine.

Appendix A. Decomposition of housing and financial wealth growth rates into frequency (persistence) components

This section describes the decomposition of the two Case et al. (2013) wealth growth rate variables (Δa_{it}^{stock} and Δa_{it}^{house}) defined in Section 3 each into the three persistence components used in Eq. 1.⁹ A more extensive

⁸ Appendix C provides more detail on the estimation results; in particular, Table C.1 lists the coefficient and standard error estimates for all of the non-wealth coefficients. Summarizing, these estimated coefficient on retail sales lagged two quarters is statistically significant; the estimated coefficient on a third lag in this variable is statistically insignificant when this variable is included in the model specification, so two lags were necessary and also sufficient. The estimated coefficients on these two lags in the dependent variable sum to less than one, as they should for a dynamically stable estimated model. Finally, the estimated coefficient on the error-correction term is statistically significant and has the negative sign one would expect.

⁹ We carry the state subscript (i) throughout this section solely for notational consistency. This section is only concerned with the persistence decomposition of the time-series data from one particular (the “ i^{th} ”) state.

description is available in Ashley and Verbrugge (2006, 2009) and Ashley et al. (2014).

The general idea of the Ashley–Verbrugge decomposition approach is to partition a time-series into frequency (persistence) components using a discrete-time Fourier transform applied to a moving window of length M periods. This transformation is embodied in the $M \times M$ orthogonal matrix A . In the present application, where M is specified to be of length sixteen quarters, the (q, j) element of the matrix A is defined as:

$$a_{s,j} = \begin{cases} \left(\frac{1}{16}\right)^{\frac{1}{2}}, & \text{for } q = 1 \\ \left(\frac{2}{16}\right)^{\frac{1}{2}} \cos\left(\frac{\pi q(j-1)}{16}\right), & \text{for } q = 2, 4, 6, \dots, 14 \\ \left(\frac{2}{16}\right)^{\frac{1}{2}} \sin\left(\frac{\pi(q-1)(j-1)}{16}\right), & \text{for } q = 3, 5, 7, \dots, 15 \\ \left(\frac{1}{16}\right)^{\frac{1}{2}} (-1)^{j+1}, & \text{for } q = 16 \end{cases} \quad (\text{A1})$$

While at first glance somewhat opaque, this transformation is actually quite intuitive once it is unpacked a bit. Imagine applying the transformation A to w_t , the sixteen quarterly observations on housing wealth ($\Delta a_{it}^{\text{house}}$) comprising a particular window of data on the time series for state i . The first row of A is just a constant, so the series Aw_t amounts to a sixteen-quarter moving average of the housing wealth data. As the window moves through the data set, this component (based on the first row of A) picks up any non-linear trend in $\Delta a_{it}^{\text{house}}$ and leads to the zero-frequency (highest persistence) component of $\Delta a_{it}^{\text{house}}$.¹⁰

The elements in the second and third rows of A vary sinusoidally, completing one complete cycle across their sixteen columns. Each of these rows (multiplied by w_t) picks up the portion of the sample variation in w_t that varies slowly across the sixteen quarters and ignores any fluctuations in w_t that essentially reverse themselves fairly quickly – i.e., within a few quarters.

In contrast, the final (sixteenth) row of the A matrix is $(\frac{1}{4}, -\frac{1}{4}, \frac{1}{4}, -\frac{1}{4}, \dots, \frac{1}{4}, -\frac{1}{4})$. This row (multiplied by w_t) picks up only the portions of the sample variation in w_t which vary quite quickly, essentially reversing themselves within a couple of quarters. Indeed, if w_t is close to constant over the sixteen quarters of data in the window, then this sixteenth component of Aw_t will be almost zero.

Similarly, the middle rows of A extract the sample variation w_t with moderate persistence – i.e., which tends to “self-reverse” within a period of a year or two.

The vector Aw_t thus has sixteen components, corresponding to the nine distinct frequencies allowed by a window of this length.¹¹ To obtain the time domain values for this window corresponding to, for example, the lowest non-zero frequency, one need only zero out the first component and components four through sixteen of Aw_t and then

left-multiply the result by the matrix A^t , which is the inverse of A . This yields a vector whose last (sixteenth) element is the period- t component of the wealth vector w_t for this window corresponding to this particular frequency. The components of w_t corresponding to the other frequencies are obtained in similar fashion, by zeroing out other components of Aw_t . Finally, the components of w_{t+1} at each frequency are obtained by moving the window ahead one period, so as to end in period $t+1$ – i.e., by using w_{t+1} instead of w_t in the procedure just described.¹²

The foregoing decomposition is substantially preferable than that which could have been obtained by simply HP filtering each wealth series, as neither the filtered series nor its deviation from the original wealth data would in that case have any specific interpretation as a component with a well-defined persistence level or reversion period. Nor, relatedly, could such a filtration method model yield components with more than two levels of persistence.¹³

The filtering (partitioning) of $\Delta a_{it}^{\text{house}}$ or $\Delta a_{it}^{\text{stock}}$ into persistence components does involve a few technical complications; see Ashley and Tsang (2012, 2014) for details on these.¹⁴ However, a couple of these complications are worth at least pointing out here.

As noted above, limiting the window length implies that any fluctuations in the time series which are so persistent as to be essentially constant over a period of this length are relegated to the lowest (zero-frequency) component, along with any (possibly non-linear) trend. In fact, with a sixteen-quarter window, fluctuations with a self-reversal period in excess of sixteen quarters are indistinguishable from this trend. On the other hand, our use of a moving window of limited length imparts several important countervailing advantages which make up for this limitation.

In particular, this moving window approach allows us to make the decomposition a strictly one-sided filtering: that is, each frequency (or persistence) component of the time series at time t depends only on data in time period t and previous periods. This is accomplished by using as the filtered components of the time series at time t only

¹² Note that the foregoing process amounts to applying a particular fixed, albeit nonlinear, filter to the original wealth data, a filter which is completely determined by the window length choice. No estimation is involved, so there is no reason to worry about ‘generated variable’ bias in the estimated standard errors for coefficients obtained by replacing the original wealth data by these frequency components in the regression model.

¹³ One could, on the other hand, reasonably contemplate applying a standard bandpass filter – i.e., Baxter and King (1999) – repeatedly to the wealth data in each window, to produce components identified by frequency which still add up to the original series. Such filters can even be optimal for arbitrarily long windows, albeit not for the relatively short windows used here. But we find these filters to be opaque – i.e., essentially a ‘black box’ – for anyone not an expert in spectral analysis. Nevertheless, because this approach is feasible, it was tried out in Ashley et al. (2014); there it was found to yield regression inference results very similar to those obtained using the decomposition procedure described here.

¹⁴ In particular, the discussion above suppresses the fact that Fourier analysis with windows this short requires that the window using data up through period t must be extended by several periods, using simple forecasts of the data for the following periods. This complication eliminates undesirable “end-effects” in the decomposition and is pre-programmed into the implementing computer codes.

¹⁰ Were the data so persistent as to be $I(1)$, the zero-frequency component would also include a linear time trend estimated to the data for each window.

¹¹ Looking at the rows of A as defined in Eq. A1, the first row corresponds to a frequency of zero – or fluctuations of arbitrarily large persistence. Rows two and three correspond to the lowest non-zero frequency – or fluctuations with a self-reversal period of sixteen quarters; rows four and five both correspond to fluctuations with a period of $\frac{16}{2}$ or eight quarters. And so forth, with the sixteenth row corresponding to the ninth (highest) frequency and a self-reversal period of just two quarters.

the last (period t) value obtainable from the window ending in period t , then – as noted above – moving the window along one period, so as to obtain the filtered components for period $t + 1$, and so forth. This feature is essential because Fourier transformation inherently mixes future and past values together. Thus, if all of the data were dealt with at once, rather than using this moving window approach, then OLS estimation of regression model parameters using these frequency components would be inconsistent if the original time series is in a feedback relationship with the dependent variable. The dependent variable here is the growth rate of state-level retail sales, so it is not unreasonable to fear feedback between it and the wealth growth rate variables. Thus, one-sided filtering is essential.

Finally – per the discussion above – windows limited to sixteen quarters in length necessarily restrict the number of possible distinct frequencies to just nine. Thus, each wealth time series can be partitioned into nine frequency (persistence) levels which add up, by construction, to the original wealth time series. We could have simply replaced each of the two wealth series in the regression model by these nine frequency components, but we instead chose to aggregate these nine components into the three components discussed in Section 1: a “low-frequency” component corresponding to the nonlinear (moving average) trend (including any fluctuations with a reversion period in excess of sixteen quarters); a “mid-frequency” component corresponding to any fluctuations with a reversion period more than four, but less than sixteen quarters in length; and a “high-frequency” component corresponding to all fluctuations with reversion period of four quarters or less. This aggregation reduces the number of regression coefficients to be estimated – from sixteen additional coefficients for the two decomposed wealth variables down to just four additional coefficients. The big payoff from it, however, is that these aggregated components are more economically interpretable, as corresponding to fluctuations which persist for more than four years, for between one and four years, and for a year or less, respectively.

Appendix B. Integration and Co-integration issues

As one might expect from the discussion in Section 1, we find that the levels variables (s_{it} , a_{it}^{house} , a_{it}^{stock} , and y_{it}) are integrated of order one – i.e., $I(1)$ – for each state. Each of these variables is therefore modeled, as in Case et al. (2013), in growth rate form; we then tested for co-integration.

We first conducted the Sarno (2000) multivariate augmented Dickey-Fuller (*MADF*) panel unit root test. The null hypothesis that the data for all fifty-one states are non-cointegrated is rejected with a P -value less than 0.0005. We then conducted the Johansen and Juselius (1990) test for the data on each state, so as to investigate the number of cointegrating relationships in each case. This testing indicates that there is at most a single cointegrating relationship in roughly half of the states and none in the rest.

Table B.1
ADF cointegration test results.

State	Test statistic	P -value for H_0
Alabama	−3.147	0.023
Arizona	−3.804	0.003
California	−3.138	0.024
Colorado	−3.570	0.006
D.C.	−3.096	0.027
Idaho	−3.158	0.023
Illinois	−2.926	0.042
Indiana	−3.280	0.016
Iowa	−4.369	0.000
Kansas	−4.416	0.000
Kentucky	−3.195	0.020
Louisiana	−3.963	0.002
Minnesota	−3.434	0.010
Mississippi	−4.880	0.000
Missouri	−3.830	0.003
Montana	−4.219	0.001
Nebraska	−3.571	0.006
New Mexico	−3.679	0.004
North Carolina	−2.864	0.050
Oklahoma	−3.412	0.011
Pennsylvania	−3.694	0.004
South Carolina	−3.319	0.014
Utah	−4.071	0.001
Washington	−3.927	0.002
Wisconsin	−3.247	0.017
Wyoming	−2.998	0.035

Note: The null hypothesis tested here is that the error correction term is $I(1)$ – i.e., there is not a valid co-integration relationship for this state. Thus, ECM_{it} is a valid error-correction term, at the 5% level of significance, for each state listed here.

We estimate a long run co-integrating relationship between retail sales, income and wealth in each of these states, of form¹⁵:

$$ECM_{it} = s_{it} - \hat{\lambda}_1 a_{it}^{stock} - \hat{\lambda}_2 a_{it}^{house} - \hat{\lambda}_3 y_{it} \quad (B1)$$

Specifically, we estimated Eq. B1 for each state, and then ran an *ADF* unit root test on the resulting fitting errors. If the null hypothesis of $I(1)$ is rejected at the 5% level, then Eq. B1 is a valid co-integrating relationship for that state; Table B.1 lists the twenty-six states for which this is the case. We then defined the error-correction term to have value zero for all observations in the states for which no valid co-integration relationship was found and to equal ECM_{it} for the others. We find that the error correction term, so defined, enters the estimated fixed-effects regression model, lagged, with a statistically significant (and negative) coefficient: $P < 0.0005$ on the two-tailed test.

Appendix C. Model estimation details

Following Case et al. (2005, 2013), we consider it highly likely that income and both wealth variables are endogenous in Eq. 1: there is no compelling reason to think that partitioning either of the two wealth variables into the three frequency components discussed above eliminates this endogeneity. Therefore, Eq. 1 was estimated as a

¹⁵ Case et al. (2013) specify an error correction term $e_{it} = c_{it-1} - y_{it-1}$ based on their rejection of a unit root in this time series in their 2005 paper, but do not include it in most of their models.

dynamic panel data model via 2-step GMM¹⁶ using three and four lags of the growth rate in retail sales, and using two, three, and four lags of the growth rate in income, $\Delta a_{it}^{\text{wealth_type_low}}$, $\Delta a_{it}^{\text{wealth_type_mid}}$ and $\Delta a_{it}^{\text{wealth_type_high}}$ as instruments for the growth rate in retail sales, the growth rate in income, $\Delta a_{it}^{\text{wealth_type_low}}$, $\Delta a_{it}^{\text{wealth_type_mid}}$ and $\Delta a_{it}^{\text{wealth_type_high}}$. Two lags of the growth rate in retail sales (the dependent variable) were needed as explanatory variables, so the second lag of this variable was not available as an instrument.¹⁷

The asymptotic validity of all of the standard error estimates and hypothesis test *P*-values discussed in the tables below hinges on an assumption that the model error (e_{it}) in Eq. 1 is well-behaved. Serial non-autocorrelation in e_{it} is ensured by the inclusion of a sufficient number of lags in the dependent variable; two-step GMM estimation ensures that any heteroskedasticity in e_{it} is not problematic.

However, plots of the fitting errors for each of the 51 states display three isolated outliers amongst the 6,885 sample observations, arising from 135 quarters of data for each of the 51 states. These outliers in the fitting errors arise because the Case et al. (2013) data set contains three notable anomalies, which consist of very large values for components of (Δs_{it} , $\Delta a_{it}^{\text{stock}}$, $\Delta a_{it}^{\text{house}}$) in 1998Q4 for Idaho, in 1998Q4 for Illinois, and in 2003Q4 for Maryland. The model was consequently re-estimated with these three states omitted from the sample; this re-estimation yielded very similar coefficient estimates, but (for the lowest and the highest frequency bands) notably more precise inferences on the key results regarding the larger impact of housing wealth (versus stock wealth) fluctuations on retail sales. Therefore, all of the results discussed in Section 4 are based on the model dropping the data for these three states.¹⁸

Finally, because of the relevance and intensity of the financial crisis beginning in 2008, an explicit allowance was made for the possibility of a structural shift in all of the estimated model coefficients, beginning in the first quarter of 2008.¹⁹ This was accomplished by defining two

¹⁶ See Baum et al. (2007) and Arellano and Bond (1991); their estimation procedure is implemented in the Stata command *ivreg2* (*gmm2s*).

¹⁷ Using two lags (and more) of the dependent variable and of the other explanatory variables as instruments is the standard procedure in dynamic panel data models like this one. Specifically, Judson and Owen (1999) find that a 'restricted GMM' estimator that uses a subset of the available lagged values as instruments increases computational efficiency without materially reducing statistical performance. Starting from the second lag of these instruments helps meet several potential objections to the IV model, as discussed in Campbell and Mankiw (1990).

¹⁸ With a total of 6,885 observations, the omission of the data from these three states (comprising just 6% of the total sample) creates only a trivial loss in the actual estimation precision. (Recall that estimation precision varies, roughly, as the reciprocal of the square root of the sample size; so 6% more data yields only 3% more precision.) The displayed precision in the estimated coefficients is, of course, actually enhanced by dropping the data from these three states, because the estimated variance of the model error term is no longer upwardly biased by the presence of the outliers. Our data set is now a sample – rather than a complete collection – of the states, albeit hopefully a reasonably representative one. This seems preferable to including questionable data series for these three states and then having to use dummy variables in the estimation model (to allow for the anomaly in the retail sales data for Maryland) and interpolation – prior to the frequency decomposition – to artificially eliminate the anomalies in the stock wealth data for Idaho and Illinois.

¹⁹ Our results are insensitive to variations in the starting period for this dummy variable.

Table C.2

Estimation results for Eq. 1: coefficients specific to 2008Q1–2012Q4 sub-period.

	Lowfreq	Midfreq	Highfreq	Lowfreq	Midfreq	Highfreq
	Stock	Stock	Stock	House	House	House
	0.0467	0.292***	−0.299***	−0.255***	0.110***	−0.145***
	(0.0458)	(0.0570)	(0.0332)	(0.0644)	(0.0127)	(0.0249)

Note: ****P* < 0.01, ***P* < 0.05, **P* < 0.1. Robust standard errors in parentheses.

dummy variables: one which is set to one for every period from 1975Q1 through 2007Q4 (and zero thereafter) and another which is set to one only for the period commencing with quarter 2008Q1. Interacting each of these two dummy variables with all of the model explanatory variables made it convenient to separately estimate all of the coefficient estimates (along with their concomitant standard error estimates) for both periods. As noted in Section 4, the joint null hypothesis that each coefficient is the same for both periods is rejected with *P* < 0.0005. Thus there is very strong evidence of a structural shift at this point. Unfortunately, the resulting parameter estimates for this later period, displayed in Table C.2 are quite wild, with significant and perverse signs on several wealth variables. We conclude, therefore, that the Eq. 1 model specification of Case et al. (2013) simply breaks down at this point and decline to interpret the coefficient estimates for this latter period.

References

- Arellano, M., Bond, S., 1991. Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *Rev. Econ. Stud.* 58 (2), 277–297.
- Ashley, R., Tsang, K.P., 2012. The oil price-macroeconomy relationship: Does persistence matter? Available at SSRN 2185548.
- Ashley, R., Verbrugge, R.J., 2006. Mis-specification in Phillips curve regressions: quantifying frequency dependence in this relationship while allowing for feedback. Technical report.
- Ashley, R., Verbrugge, R.J., 2009. Frequency dependence in regression model coefficients: an alternative approach for modeling nonlinear dynamic relationships in time series. *Econ. Rev.* 28 (1–3), 4–20.
- Ashley, R., Tsang, K.P., Verbrugge, R., 2014. Frequency dependence in a real-time monetary policy rule. Available at SSRN 1543928.
- Baum, C.F., Schaffer, M.E., Stillman, S., 2007. Enhanced routines for instrumental variables/GMM estimation and testing. *Stata J.* 7 (4), 465–506.
- Baxter, M., King, R.G., 1999. Measuring business cycles: approximate band-pass filters for economic time series. *Rev. Econ. Stat.* 81 (4), 575–593.
- Belsky, E.S., 2010. Housing wealth effects and course of the US economy: Theory, evidence, and policy implications, *The Blackwell Companion to the Economics of Housing: The Housing Wealth of Nations* pp. 82–104.
- Benjamin, J.D., Chinloy, P., Jud, G.D., 2004. Real estate versus financial wealth in consumption. *J. Real Estate Finance Econ.* 29 (3), 341–354.
- Campbell, J.Y., Mankiw, N.G., 1990. Permanent income, current income, and consumption. *J. Bus. Econ. Stat.* 8 (3), 265–279.
- Carroll, C.D., Otsuka, M., Slacalek, J., 2011. How large are housing and financial wealth effects? a new approach. *J. Money Credit Banking* 43 (1), 55–79.
- Case, K.E., Quigley, J.M., Shiller, R.J., et al., 2005. Comparing wealth effects: the stock market versus the housing market. *B.E. J. Macroecon.* 5 (1), 1.
- Case, K.E., Quigley, J.M., Shiller, R.J., 2013. Wealth effects revisited: 1975–2012. *Crit. Finance Rev.* 2 (1), 101–128.
- Dvornak, N., Kohler, M., 2007. Housing wealth, stock market wealth and consumption: a panel analysis for Australia. *Econ. Rec.* 83 (261), 117–130.
- Edison, H., Silk, T., 2001. Wealth effects and the new economy, IMF Working Paper.

- Elbourne, A., 2008. The UK housing market and the monetary policy transmission mechanism: an SVAR approach. *J. Housing Econ.* 17 (1), 65–87.
- Engelhardt, G.V., 1996. House prices and home owner saving behavior. *Reg. Sci. Urban Econ.* 26 (3), 313–336.
- Johansen, S., Juselius, K., 1990. Maximum likelihood estimation and inference on cointegration with applications to the demand for money. *Oxford Bull. Econ. Stat.* 52 (2), 169–210.
- Judson, R.A., Owen, A.L., 1999. Estimating dynamic panel data models: a guide for macroeconomists. *Econ. Lett.* 65 (1), 9–15.
- Sarno, L., 2000. Real exchange rate behaviour in high inflation countries: empirical evidence from Turkey, 1980–1997. *Appl. Econ. Lett.* 7 (5), 285–291.
- Shirvani, H., Wilbratte, B., 2011. The wealth effect of the stock market revisited. *J. Appl. Bus. Res. (JABR)* 18 (2).